WHAT IS CLAIMED IS:

1. A system for providing multi-path communications in a virtualization storage environment for managing a virtual volume of data, comprising:

a host system connected to first and second communication fabrics;

a network switch system connected to the first and second communication fabrics and to third and fourth communication fabrics; and

a set of storage devices storing virtual volume data and connected to the third and fourth storage communication fabrics,

wherein the network switch system includes:

a first set of storage processors having access to the host system and maintaining virtual volume objects reflecting a logical configuration of the virtual volume,

a second set of storage processors having access to the storage devices and maintaining virtual volume objects associated with logical partitions of the virtual volume data,

wherein the first and second set of storage processors are interconnected by a fifth communication fabric and the network switch system provides fault tolerant access by the host system to the virtual volume data using one of a plurality of dynamically configurable multi-communication paths traversing selective combinations of the fabrics, storage processors, and storage devices.

2. The system of claim 1, wherein the network switch system dynamically configures a multi-communication path to provide access to the virtual volume data based on the availability of at least one of the first through fifth communication fabrics.

3. The system of claim 1, wherein the network switch system dynamically configures a multi-communication path to provide access to the virtual volume data based on the availability of at least one of the first and second sets of storage processors.

4. The system of claim 1, wherein the network switch system further includes:

a first blade component including a first subset of the first set storage processors; and

a second blade component including a second subset of the first set storage processors,

wherein the first blade component is attached to the first communication fabric and the fifth communication fabric and the second blade component is attached to the second communication fabric and the fifth communication fabric.

5. The system of claim 4, wherein when the second blade component is inaccessible by the host system the network switch system dynamically configures a first multi-communication path including the host system, the second communication fabric, and the first blade component, and

when the first blade component is inaccessible by the host system the network switch system dynamically configures a second multi-communication path including the host system, the first communication fabric, and the first blade component.

6. The system of claim 5, wherein the fifth communication fabric is a redundant communication fabric including a sixth and seventh communication fabrics, and wherein the network switch system dynamically configures either one of the first and second multi-communication paths to include the sixth communication fabric when the fifth communication fabric is unavailable.

7. The system of claim 5, wherein the fifth communication fabric is a redundant communication fabric including a sixth and seventh communication fabrics, and wherein the network switch system dynamically configures either one of the first and second multi-communication paths to include the seventh communication fabric when the sixth communication fabric is unavailable.

8. The system of claim 4, further including:

a third blade component including a first subset of the second set storage processors; and

a fourth blade component including a second subset of the second set storage processors,

wherein the third blade component is attached to the third communication fabric and the fifth communication fabric and the fourth blade component is attached to the fourth communication fabric and the fifth communication fabric.

9. The system of claim 5, further including:

a third blade component including a first subset of the second set storage processors; and

a fourth blade component including a second subset of the second set storage processors,

wherein the third blade component is attached to the third communication fabric and the fifth communication fabric and the fourth blade component is attached to the fourth communication fabric and the fifth communication fabric.

10. The system of claim 9, wherein when the third blade component is inaccessible, the network switch system dynamically configures a first multi-communication path including at least the host system, at least one of the first and second communication fabrics, at least one of the first and second blade components, the fifth communication fabric, and the fourth blade component, and

when the fourth blade component is inaccessible, the network switch system dynamically configures a second multi-communication path including the at least the host system, at least one of the first and second communication fabrics, at least one of the first and second blade components, the fifth communication fabric, and the third blade component.

11. The system of claim 10, wherein the fifth communication fabric is a redundant communication fabric including a sixth and seventh communication fabrics, and wherein the network switch system dynamically configures either one of the first and second multi-communication paths to include the sixth communication fabric when the fifth communication fabric is unavailable.

12. The system of claim 10, wherein the fifth communication fabric is a redundant communication fabric including a sixth and seventh communication fabrics, and wherein the network switch system dynamically configures either one of the first and second multi-communication paths to include the seventh communication fabric when the sixth communication fabric is unavailable.

13. The system of claim 1, wherein each of the storage devices includes a first port connected to the third communication fabric and a second port connected to the fourth communication fabric and the network switch system dynamically configures a multi-communication path to include either one of the first and second ports of a given storage device when the given storage device includes virtual volume data referenced by the host system.

14. The system of claim 13, wherein the network switch system dynamically configures a multi-communication path to include the first port of the given storage device when the second port is unavailable and configures the multi-communication

path to include the second port of the given storage device when the first port is unavailable.

15. The system of claim 8, wherein each of the storage devices includes a first port connected to the third communication fabric and a second port connected to the fourth communication fabric and the network switch system dynamically configures a multi-communication path to include either one of the first and second ports of a given storage device when the given storage device includes virtual volume data referenced by the host system.

16. The system of claim 15, wherein the network switch system dynamically configures the multi-communication path to include the first port of the given storage device when the fourth communication fabric is unavailable and configures the multi-communication path to include the second port of the given storage device when the third communication fabric is unavailable.

17. The system of claim 9, wherein each of the storage devices includes a first port connected to the third communication fabric and a second port connected to the fourth communication fabric and the network switch system dynamically configures a multi-communication path to include either one of the first and second ports of a given storage device when the given storage device includes virtual volume data referenced by the host system.

18. The system of claim 17, wherein the network switch system dynamically configures the multi-communication path to include the first port of the given storage device when the fourth communication fabric is unavailable and configures the multi-communication path to include the second port of the given storage device when the third communication fabric is unavailable.

19. The system of claim 1, wherein the network switch system includes:

a first blade component including a first subset of the first set storage processors;

a second blade component including a second subset of the first set storage processors;

a third blade component including a first subset of the second set storage processors; and

a fourth blade component including a second subset of the second set storage processors,

wherein the fifth communication fabric includes a sixth communication fabric and a seventh communication fabric each connected to the first through fourth blade components, and wherein a first storage device includes a first port connected to the third communication fabric and a second port connected to the fourth communication fabric.

20. The system of claim 19, wherein when the seventh communication fabric and second port are unavailable, the network switch system dynamically configures a multi-communication path to traverse the host system, any one of the first and second

communication fabrics, any one of the first and second blade components, the sixth communication fabric, any one of the third and fourth blade components, the third communication fabric, and the first port.

21. The system of claim 19, wherein when the sixth communication fabric and second port are unavailable, the network switch system dynamically configures a multi-communication path to traverse the host system, any one of the first and second communication fabrics, any one of the first and second blade components, the seventh communication fabric, any one of the third and fourth blade components, the third communication fabric, and the first port.

22. The system of claim 19, wherein when the seventh communication fabric and first port are unavailable, the network switch system dynamically configures a multi-communication path to traverse the host system, any one of the first and second communication fabrics, any one of the first and second blade components, the sixth communication fabric, any one of the third and fourth blade components, the third communication fabric, and the second port.

23. The system of claim 19, wherein when the sixth communication fabric and first port are unavailable, the network switch system dynamically configures a multi-communication path to traverse the host system, any one of the first and second communication fabrics, any one of the first and second blade components, the seventh

communication fabric, any one of the third and fourth blade components, the third communication fabric, and the second port.

24. A method for providing multi-path communications in a virtualization environment for managing a virtual volume of objects including a host system connected to a network switch system by first and second communication fabrics, and a set of storage devices storing the virtual volume data and connected to the network storage system by third and fourth communication fabrics, wherein the network switch system includes a first set of storage processors having access to the storage devices and maintaining virtual volume objects associated with logical partitions of the virtual volume data and a second set of storage processors having access to the host system and maintaining virtual volume objects reflecting a logical configuration of the virtual volume, wherein the first and second sets of storage processors are interconnected by a fifth communication fabric, the method comprising:

receiving a request from the host system to access the virtual volume data;

determining a multi-communication path that traverses selective ones of the fabrics, storage processors, and storage devices based on a current availability of at least one of the fabrics, storage processors, and storage devices; and

providing access to the requested virtual volume data over the multi-communication path.

25. The method of claim 24, wherein determining a multi-communication path includes:

dynamically configuring the multi-communication path to provide access to the virtual volume data based on the availability of at least one of the first through fifth communication fabrics.

26. The method of claim 24, wherein determining the multi-communication path includes:

dynamically configuring the multi-communication path to provide access to the virtual volume data based on the availability of at least one of the first and second sets of storage processors.

27. The method of claim 24, wherein the network switch system further includes a first blade component including a first subset of the first set storage processors and a second blade component including a second subset of the first set storage processors, and wherein the wherein the first blade component is attached to the first communication fabric and the fifth communication fabric and the second blade component is attached to the second communication fabric and the fifth communication fabric, wherein the method further includes:

dynamically configuring the first multi-communication path to traverse the host system, the second communication fabric, and the first blade component when the second blade component is inaccessible by the host system.

28. The method of claim 27, further including:

dynamically configuring the multi-communication path to traverse the host system, the first communication fabric, and the second blade component when the first blade component is inaccessible by the host system.

29. The method of claim 28, wherein the fifth communication fabric is a redundant communication fabric including a sixth and seventh communication fabrics, and wherein the method further includes:

dynamically configuring the multi-communication path to include the sixth communication fabric when the seventh communication fabric is unavailable.

30. The method of claim 28, wherein the fifth communication fabric is a redundant communication fabric including a sixth and seventh communication fabrics, wherein the method further includes:

dynamically configuring either the multi-communication path to traverse the seventh communication fabric when the sixth communication fabric is unavailable.

31. The method of claim 27, wherein the network switch system further includes a third blade component including a first subset of the second set storage processors and a fourth blade component including a second subset of the second set storage processors, wherein the third blade component is attached to the third communication fabric and the fifth communication fabric and the fourth blade component is attached to the fourth communication fabric and the fifth communication fabric, and wherein the method further includes:

dynamically configuring the multi-communication path to traverse at least the host system, at least one of the first and second communication fabrics, at least one of the first and second blade components, the fifth communication fabric, and the fourth blade component, when the third blade component is inaccessible.

32. The method of claim 31, wherein the method further includes:

dynamically configuring the multi-communication path to traverse at least the host system, at least one of the first and second communication fabrics, at least one of the first and second blade components, the fifth communication fabric, and the third blade component, when the fourth blade component is inaccessible.

33. The method of claim 32, wherein the fifth communication fabric is a redundant communication fabric including a sixth and seventh communication fabrics, and wherein the method further includes:

dynamically configuring the multi-communication path to traverse the sixth communication fabric when the seventh communication fabric is unavailable.

34. The method of claim 32, wherein the fifth communication fabric is a redundant communication fabric including a sixth and seventh communication fabrics, and wherein the method further includes:

dynamically configuring the multi-communication path to traverse the seventh communication fabric when the sixth communication fabric is unavailable.

35. The method of claim 24, wherein each of the storage devices includes a first port connected to the third communication fabric and a second port connected to the fourth communication fabric and the method further includes:

dynamically configuring the multi-communication path to traverse either one of the first and second ports of a given storage device when the given storage device includes virtual volume data referenced by the host system.

36. The method of claim 35, wherein the method further includes:

dynamically configuring the multi-communication path to traverse the first port of the given storage device when the second port is unavailable; and

dynamically configuring the multi-communication path to traverse the second port of the given storage device when the first port is unavailable.

37. The method of claim 31, wherein each of the storage devices includes a first port connected to the third communication fabric and a second port connected to the fourth communication fabric and the method further includes:

dynamically configuring the multi-communication path to traverse either one of the first and second ports of a given storage device when the given storage device includes virtual volume data referenced by the host system.

38. The method of claim 37, wherein the method further includes:

dynamically configuring the multi-communication path to traverse the first port of the given storage device when the fourth communication fabric is unavailable; and

dynamically configuring the multi-communication path to traverse the second port of the given storage device when the third communication fabric is unavailable.

39. The method of claim 27, wherein each of the storage devices includes a first port connected to the third communication fabric and a second port connected to the fourth communication fabric and the method further includes:

dynamically configuring the multi-communication path to traverse either one of the first and second ports of a given storage device when the given storage device includes virtual volume data referenced by the host system.

40. The method of claim 39, further including:

dynamically configuring the multi-communication path to traverse the first port of the given storage device when the fourth communication fabric is unavailable; and

dynamically configuring the multi-communication path to traverse the second port of the given storage device when the third communication fabric is unavailable.

41. The method of claim 24, wherein the network switch system includes a first blade component including a first subset of the first set storage processors, a second blade component including a second subset of the first set storage processors, a third blade component including a first subset of the second set storage processors, and a fourth blade component including a second subset of the second set storage processors, and wherein the fifth communication fabric includes a sixth communication fabric and a seventh communication fabric each connected to the first through fourth blade components, and wherein a first storage device includes a first port connected to the third communication fabric and a second port connected to the fourth communication fabric.

42. The method of claim 41, further including:

dynamically configuring the multi-communication path to traverse the host system, any one of the first and second communication fabrics, any one of the first and second blade components, the sixth communication fabric, any one of the third and fourth blade components, the third communication fabric, and the first port when the seventh communication fabric and second port are unavailable.

43. The method of claim 41, further including:

dynamically configuring the multi-communication path to traverse the host system, any one of the first and second communication fabrics, any one of the first and second blade components, the seventh communication fabric, any one of the third and fourth blade components, the third communication fabric, and the first port when the sixth communication fabric and second port are unavailable.

44. The method of claim 41, further including:

dynamically configuring the multi-communication path to traverse the host system, any one of the first and second communication fabrics, any one of the first and second blade components, the sixth communication fabric, any one of the third and fourth blade components, the third communication fabric, and the second port when the seventh communication fabric and first port are unavailable.

45. The method of claim 41, further including:

dynamically configuring the multi-communication path to traverse the host system, any one of the first and second communication fabrics, any one of the first and second blade components, the seventh communication fabric, any one of the third and fourth blade components, the third communication fabric, and the second port when the sixth communication fabric and first port are unavailable.

46. A method for providing multi-path communications in a storage virtualization environment including a multi-tiered network switch system that manages a virtual volume for a host system connected to the network switch system over a first and second path, wherein the network switch system includes a first virtualization layer that maintains first tier objects reflecting partitions of virtual volume data distributed across a set of storage devices connected to the first virtualization layer over third and fourth paths, and a second virtualization layer that maintains second tier objects reflecting a logical configuration of the virtual volume, wherein the first and second virtualization layers are interconnected over a fifth and sixth path, the method comprising:

receiving a request from the host system to access the virtual volume data;

determining whether communications are available between at least one of:

the host system and the network switch system over the first or second paths,

the first virtualization layer and the second virtualization layer over the fifth or sixth paths, and

the first virtualization layer and the storage devices over the third or fourth paths; and

providing access to the virtual volume data based on the determination.

47. A computer-readable medium including instructions for performing, when executed by a processor, a method for providing multi-path communications in a virtualization environment for managing a virtual volume of objects including a host system connected to a network switch system by first and second communication fabrics, and a set of storage devices storing the virtual volume data and connected to the network storage system by third and fourth communication fabrics, wherein the network switch system includes a first set of storage processors having access to the host system and maintaining virtual volume objects reflecting a logical configuration of the virtual volume and a second set of storage processors having access to the storage devices and maintaining virtual volume objects associated with logical partitions of the virtual volume data, wherein the first and second sets of storage processors are interconnected by a fifth communication fabric, the method comprising:

receiving a request from the host system to access the virtual volume data;

determining a multi-communication path that traverses selective ones of the fabrics, storage processors, and storage devices based on a current availability of at least one of the fabrics, storage processors, and storage devices; and

providing access to the requested virtual volume data over the multi-communication path.

48. A system for providing multi-path communications in a multi-tier storage virtualization environment for managing a virtual volume, the environment including a host system and a network switch system having a first virtualization layer and a second virtualization layer, wherein the first virtualization layer is redundantly connected to a set of storage devices storing virtual volume data and the second virtualization layer is redundantly connected to the host system and the first virtualization layer, the system comprising:

means for receiving a request from the host system to access the virtual volume data;

means for determining a multi-communication path that traverses a selective path including the redundant connections between the host system, the first virtualization layer, the second virtualization layer, and the set of storage devices; and

means for accessing the requested virtual volume data over using the multi-communication path.

49. A system for dynamically updating a virtual volume in a multi-tier virtualization storage environment including a set of storage devices storing virtual volume data and connected to a first set of storage processors, and a second set of storage processors connected to a host system associated with the virtual volume, wherein the virtual volume is defined by a set of virtual volume objects associated with selected ones of the first and second set storage processors, the system comprising:

means for receiving a request from the host system to adjust the virtual volume;

means for determining which virtual volume objects are affected by the request to adjust the virtual volume data;

means for updating the virtual volume based on the affected virtual volume objects and host system request; and

means for allowing the host system to access the updated virtual volume.